

SPARSE FLUID SIMULATION IN DIRECT X

ALEX DUNN - NVIDIA - DEV. TECH.

AGENDA

- ▶ Fluid in games.
- ▶ Eulerian (grid based) fluid.
- ▶ Sparse Eulerian Fluid.
- ▶ Feature Level 11.3 Enhancements!

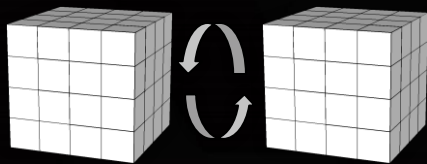
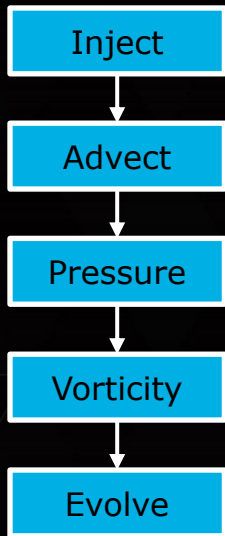
WHY DO WE NEED FLUID IN GAMES?

- ▶ Replace particle kinematics!
 - ▶ more realistic == better immersion
- ▶ Game mechanics?
 - ▶ occlusion
 - ▶ smoke grenades
 - ▶ physical interaction
 - ▶ Dispersion
 - ▶ air ventilation systems
 - ▶ poison, smoke
- ▶ Endless opportunities!

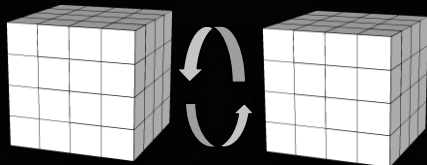


EULERIAN SIMULATION #1

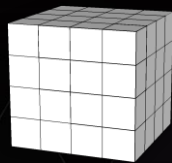
My (simple) DX11.0 eulerian fluid simulation:



2x Velocity

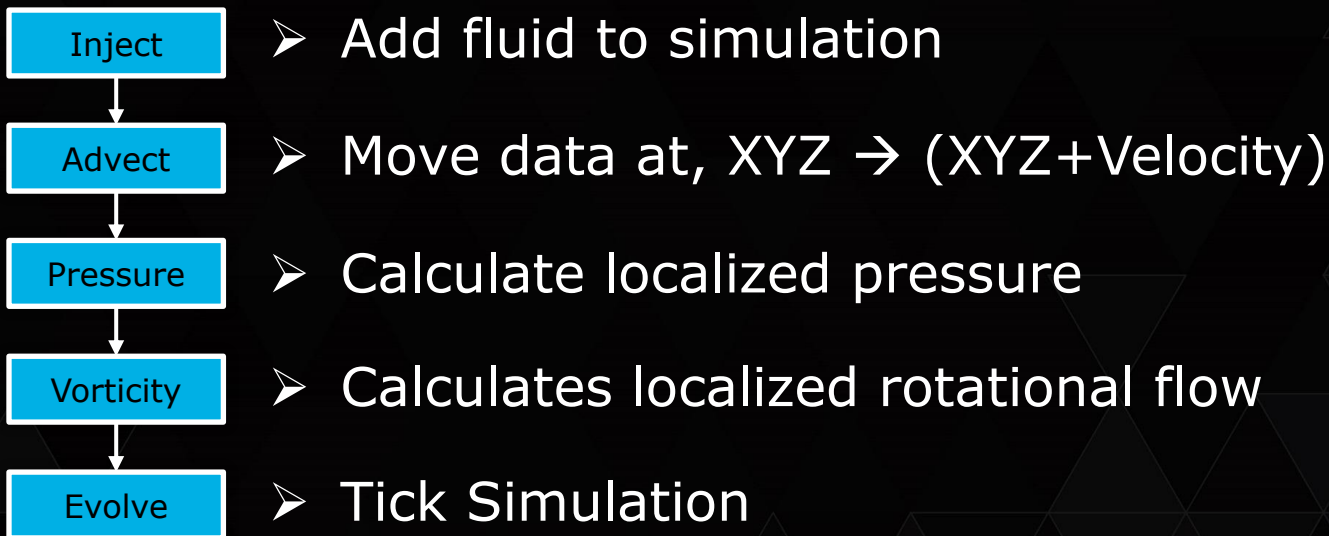


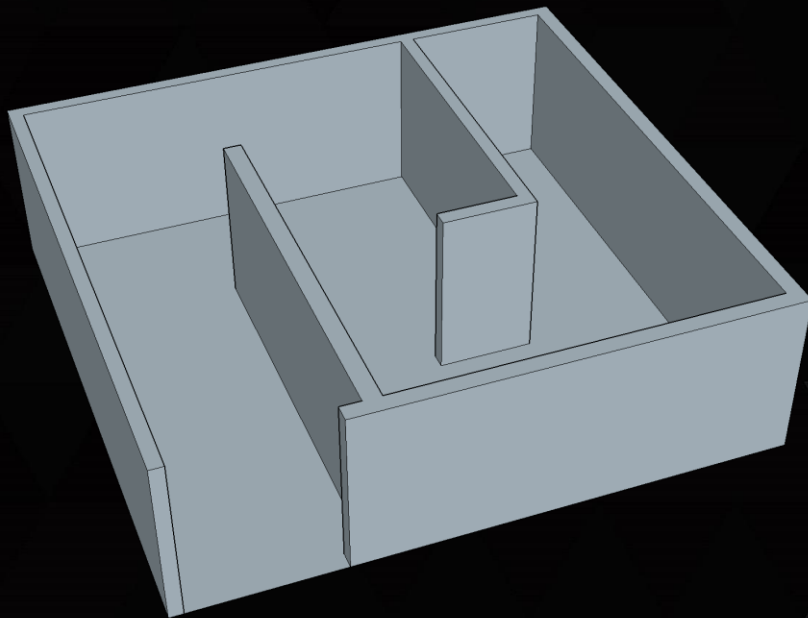
2x Pressure



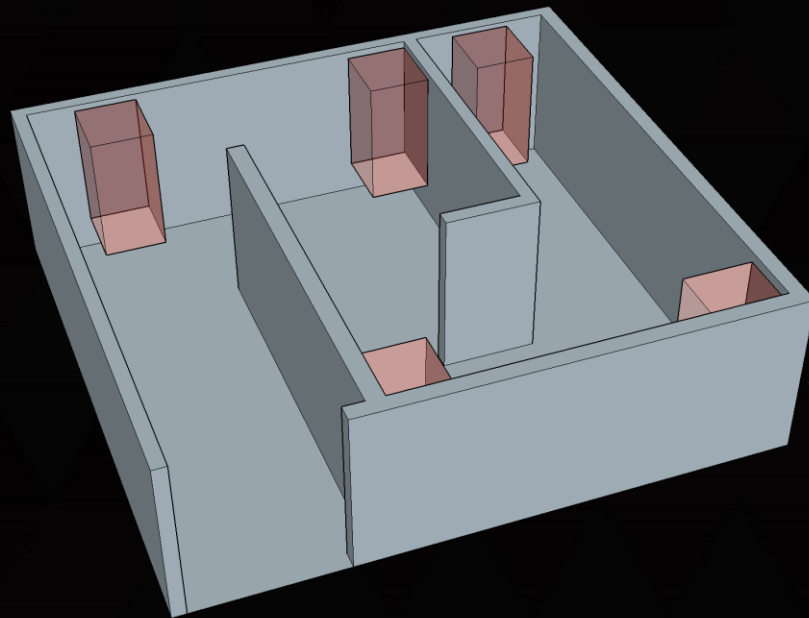
1x Vorticity

EULERIAN SIMULATION #2





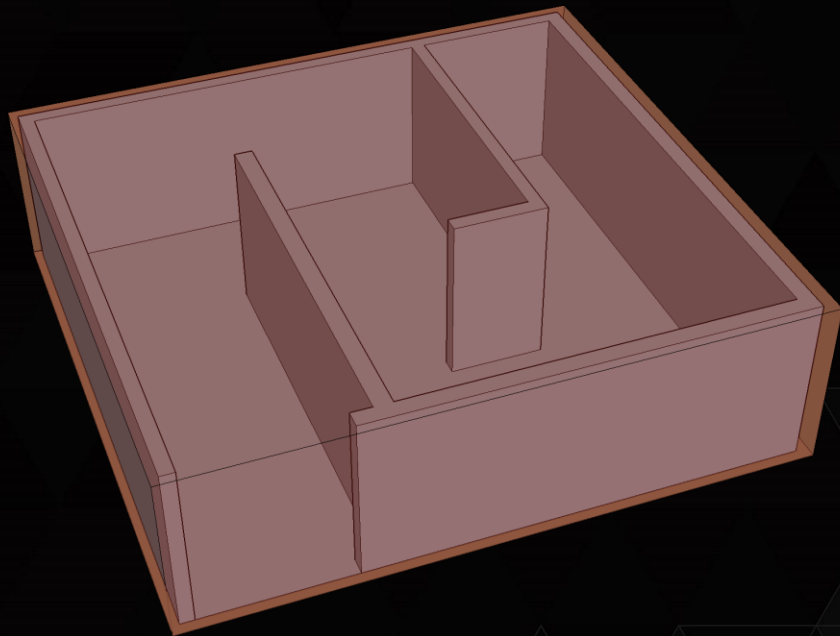
*** (some imagination required) ***



TOO MANY VOLUMES SPOIL THE...

- ▶ Fluid isn't box shaped.
 - ▶ clipping
 - ▶ wastage
- ▶ Simulated separately.
 - ▶ authoring
 - ▶ GPU state
 - ▶ no volume-to-volume interaction
- ▶ Tricky to render.

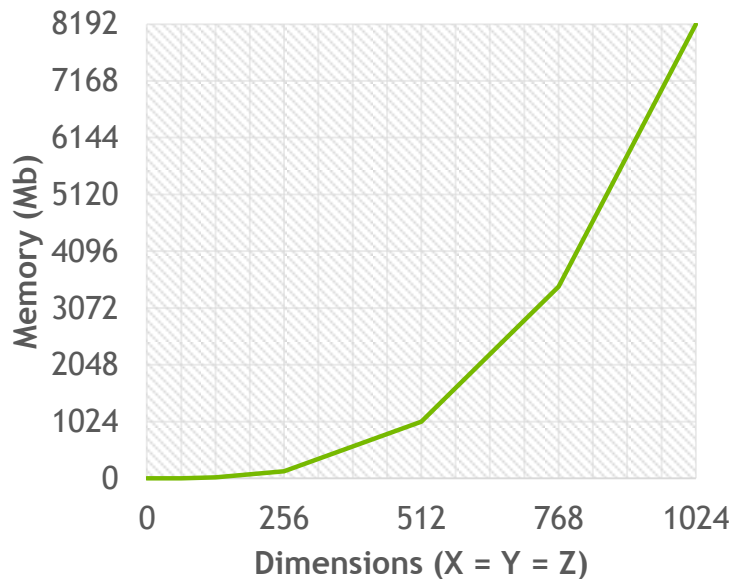




PROBLEM!

- ▶ N-order problem
 - ▶ $64^3 = \sim 0.25\text{m cells}$
 - ▶ $128^3 = \sim 2\text{m cells}$
 - ▶ $256^3 = \sim 16\text{m cells}$
 - ▶ ...
- ▶ Applies to:
 - ▶ computational complexity
 - ▶ memory requirements

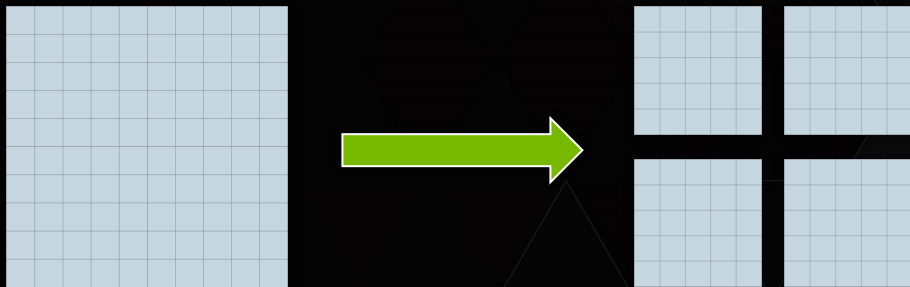
Texture3D - 4x16F



And that's just 1 texture...

BRICKS

- ▶ Split simulation space into groups of cells (each known as a brick).
- ▶ Simulate each brick independently.



BRICK MAP

- ▶ Need to track which bricks contain fluid

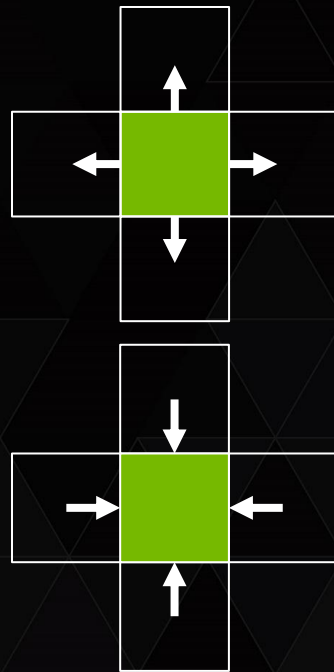
- ▶ Texture3D<uint>
- ▶ 1 voxel per brick
 - ▶ 0 → Unoccupied
 - ▶ 1 → Occupied

0	1
0	1

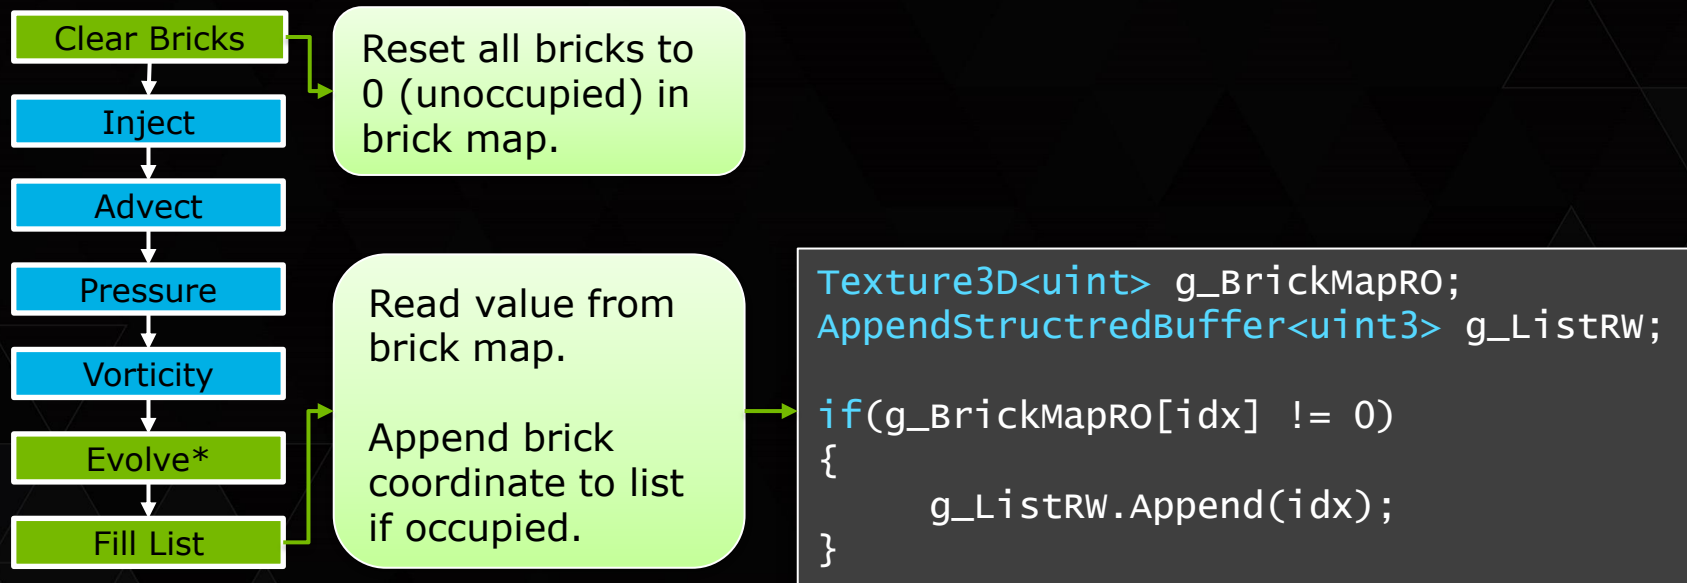
- ▶ *Could also use packed binary grids [Gruen15], but this requires atomics ☹*

TRACKING BRICKS

- ▶ Initialise with emitter
- ▶ Expansion (*unoccupied* \rightarrow *occupied*)
 - ▶ if $\{ V_{|x|y|z|} > |D_{\text{brick}}| \}$
 - ▶ expand in that axis
- ▶ Reduction (*occupied* \rightarrow *unoccupied*)
 - ▶ inverse of Expansion
 - ▶ handled automatically

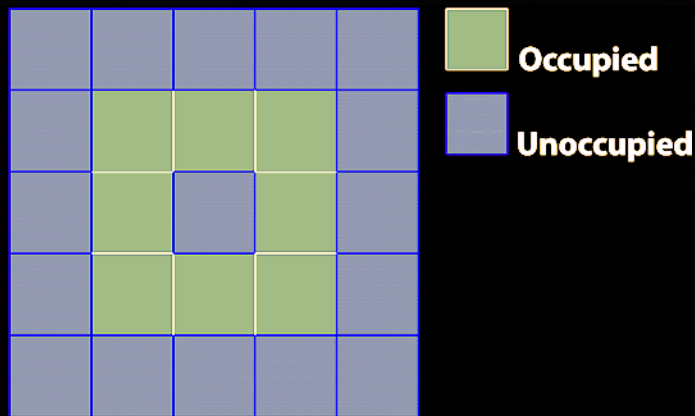


SPARSE SIMULATION



**Includes expansion*

UNCOMPRESSED STORAGE



Allocate everything; forget about unoccupied cells ☹

Pros:

- simulation is coherent in memory.
- works in DX11.0.

Cons:

- no reduction in memory usage.

COMPRESSED STORAGE

Indirection Table

	A	B	C	
	D		E	
	F	G	H	

Legend: Mapped
 Unmapped

Physical Memory

A	B	C	D	E	F	G	H
---	---	---	---	---	---	---	---

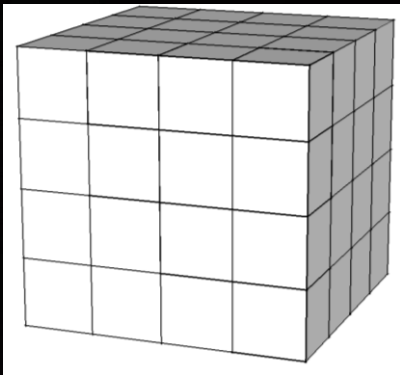
Similar to, List<Brick>

Pros:

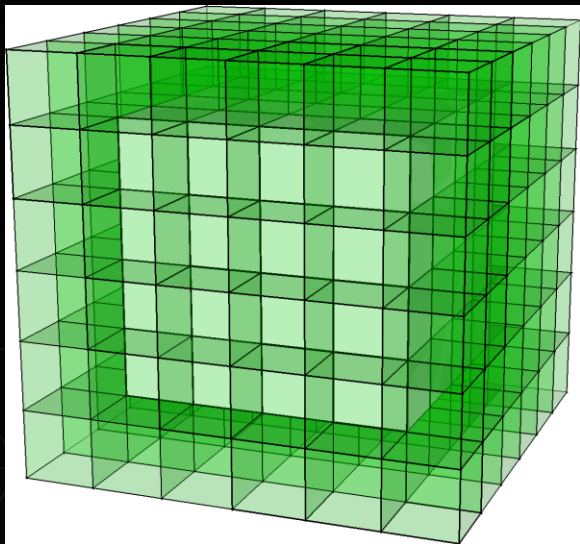
- good memory consumption.
- works in DX11.0.

Cons:

- allocation strategies.
- indirect lookup.
 - “software translation”
 - filtering particularly costly



$$1 \text{ Brick} = (4)^3 = 64$$



$$1 \text{ Brick} = (1+4+1)^3 = 216$$

- New problem;
- “ $6n^2 + 12n + 8$ ” problem.

Can we do better?

ENTER; FEATURE LEVEL 11.3

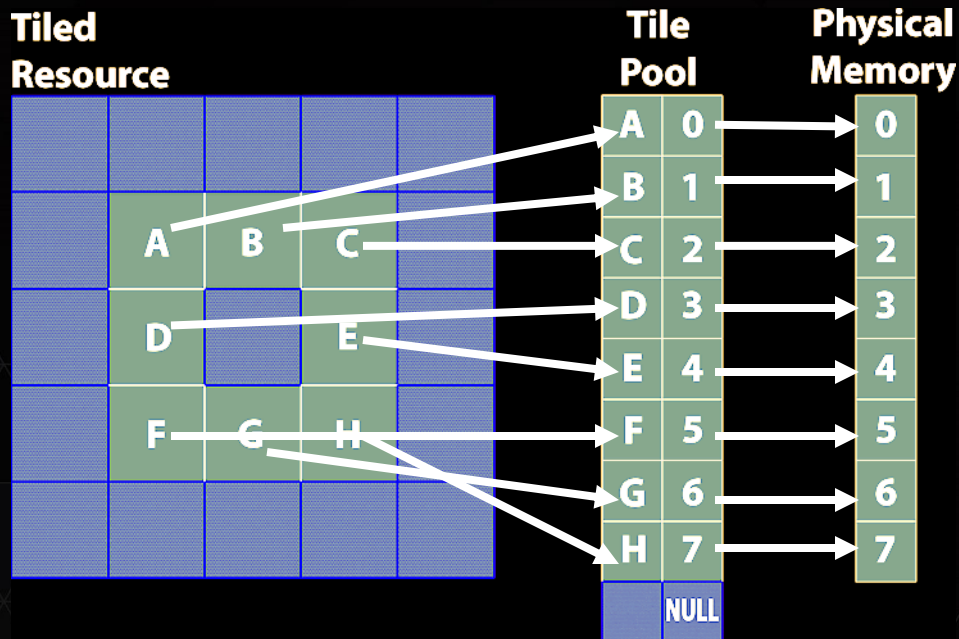
- ▶ Volume Tiled Resources (VTR)! 😊
- ▶ Extends 2D functionality in FL11.2
- ▶ Must check HW support: (DX11.3 != FL11.3)

```
ID3D11Device3* pDevice3 = nullptr;
pDevice->QueryInterface(&pDevice3);

D3D11_FEATURE_DATA_D3D11_OPTIONS2 support;
pDevice3->CheckFeatureSupport(D3D11_FEATURE_D3D11_OPTIONS2,
                             &support,
                             sizeof(support));

m_UseTiledResources = support.TiledResourcesTier ==
                     D3D11_TILED_RESOURCES_TIER_3;
```

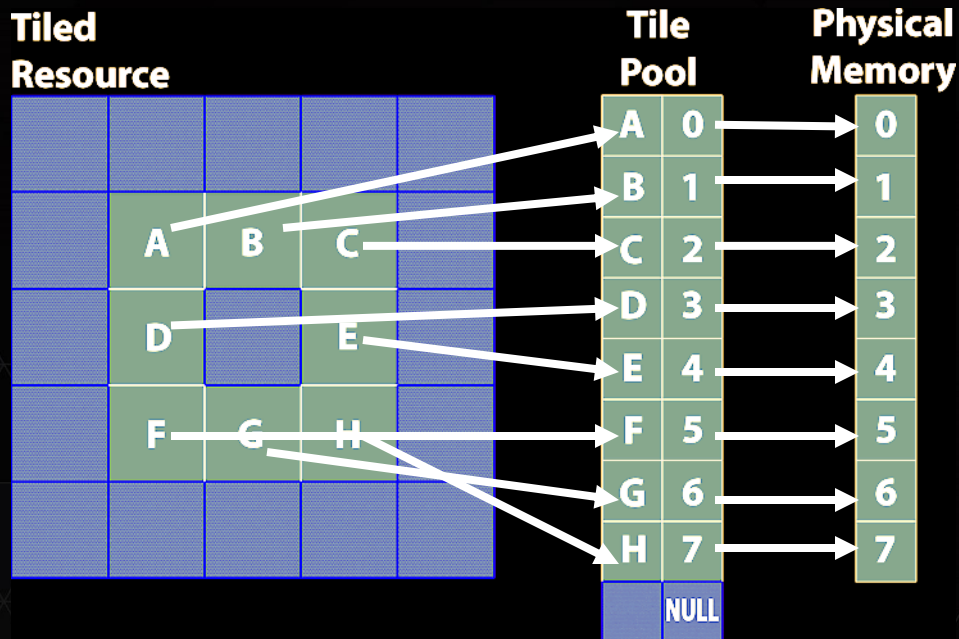
TILED RESOURCES #1



Pros:

- only mapped memory is allocated in VRAM
- “hardware translation”
- logically a volume texture
- all samplers supported
- 1 Tile = 64KB (= 1 Brick)
- fast loads

TILED RESOURCES #2



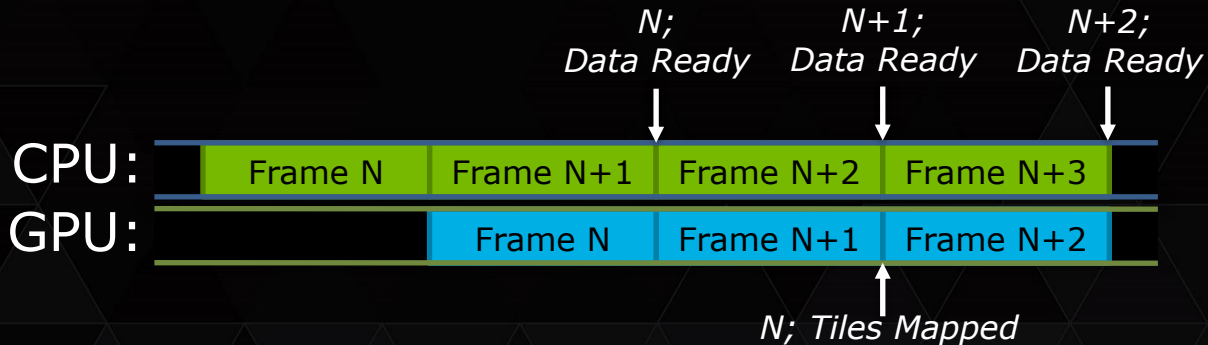
1 Tile = 64KB (= 1 Brick)

BPP	Tile Dimensions
8	64x32x32
16	32x32x32
32	32x32x16
64	32x16x16
128	16x16x16

Gotcha: Tile mappings must be updated from CPU

CPU READ-BACKS

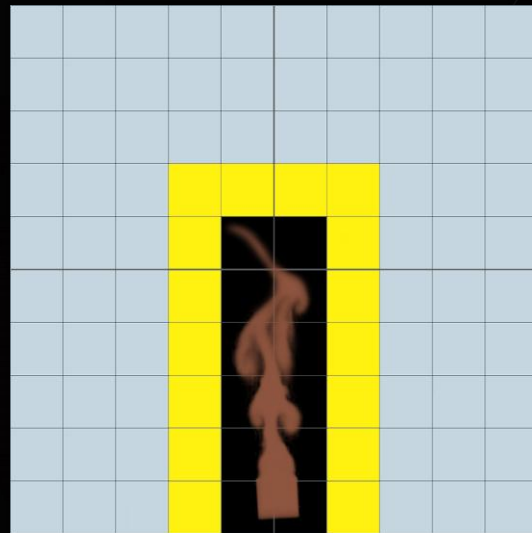
- ▶ Taboo in real time graphics
- ▶ CPU read-backs are fine, if done correctly!
 - ▶ (and bad if not)
- ▶ 2 frame latency (more for AFR in SLI)
- ▶ Profile map/unmap calls



LATENCY RESISTANT SIMULATION #1

Naïve Approach:

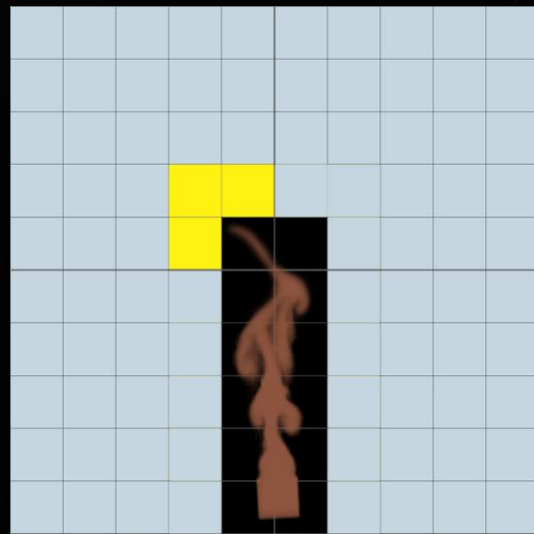
- ▶ clamp velocity to V_{\max}
- ▶ CPU Read-back:
 - ▶ occupied bricks.
 - ▶ 2 frames of latency!
- ▶ extrapolate “probable” tiles.



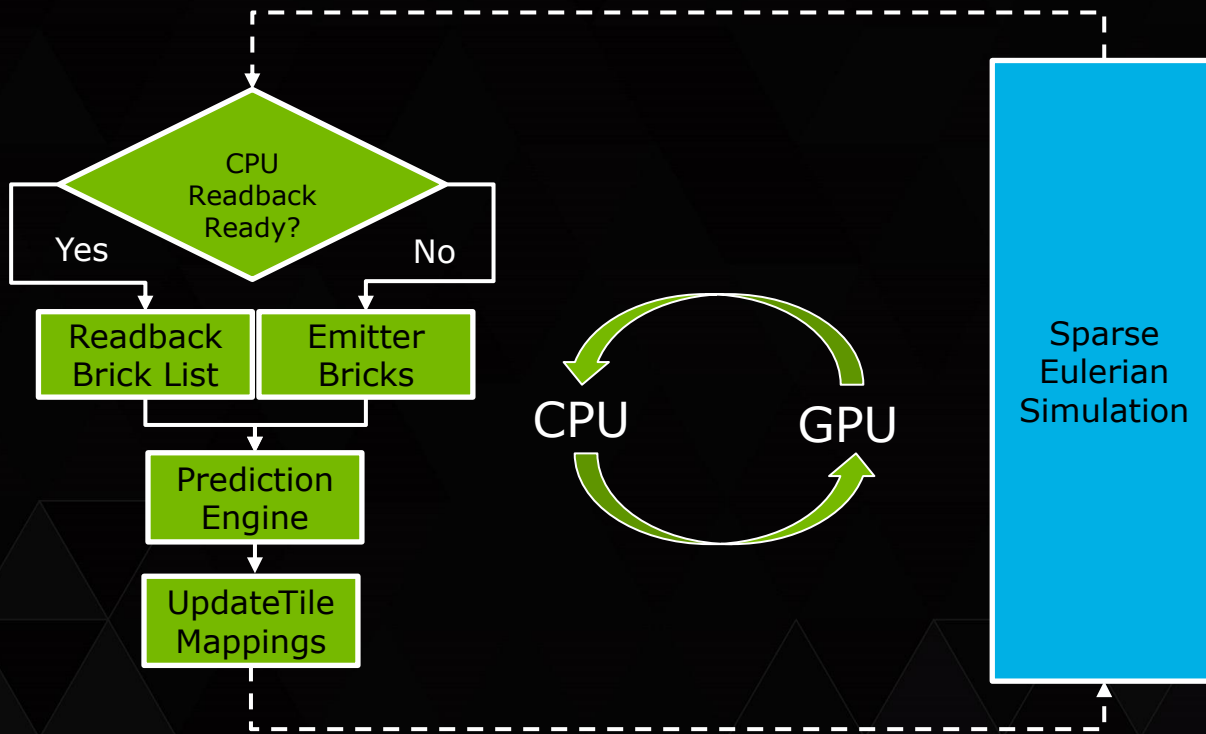
LATENCY RESISTANT SIMULATION #2

Tight Approach:

- ▶ CPU Read-back:
 - ▶ occupied bricks.
 - ▶ $\max\{|V|\}$ within brick.
 - ▶ 2 frames of latency!
- ▶ extrapolate “probable” tiles.

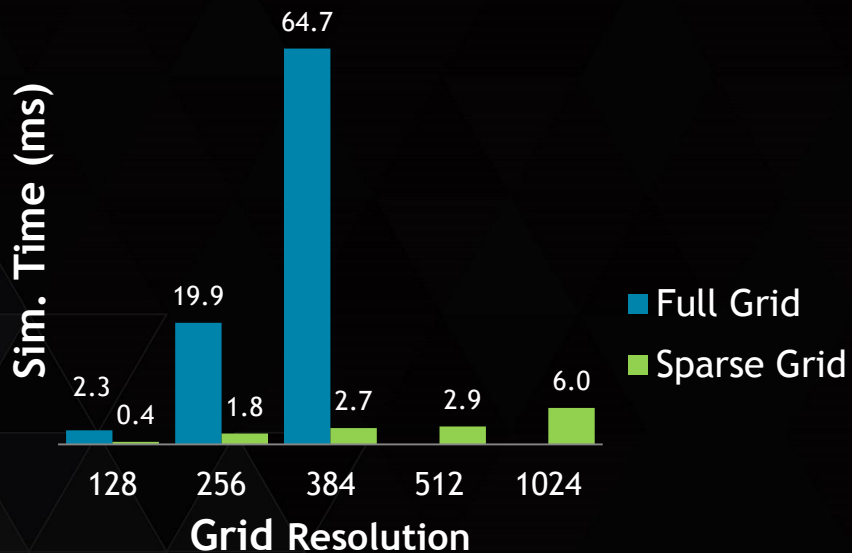


LATENCY RESISTANT SIMULATION #3



DEMO

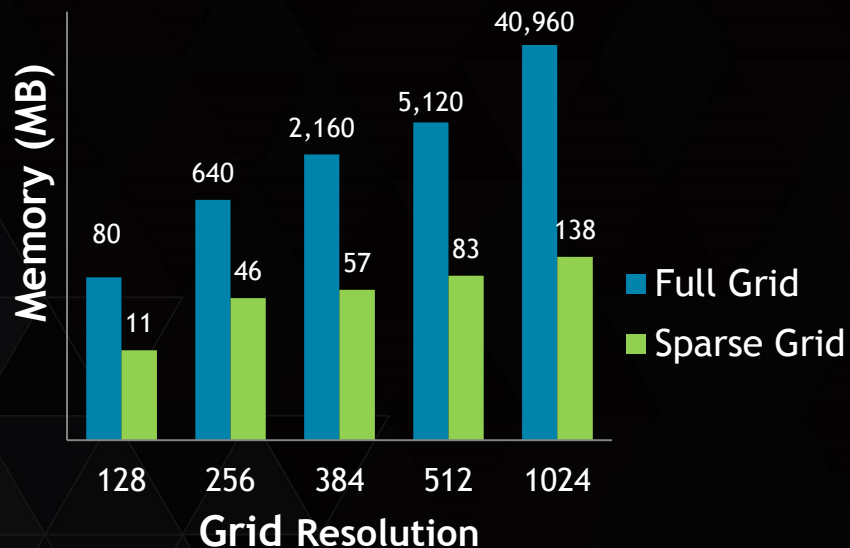
PERFORMANCE #1



	Grid Resolution				
	128 ³	256 ³	384 ³	512 ³	1,024 ³
	Full Grid				
Num. Bricks	256	2048	6,912	16,384	131,072
Memory (MB)	80	640	2,160	5,120	40,960
Simulation	2.29ms	19.04ms	64.71ms	NA	NA
	Sparse Grid				
	36	146	183	266	443
	11.25	45.63	57.19	83.13	138.44
Simulation	0.41ms	1.78ms	2.67ms	2.94ms	5.99ms
Scaling Sim.	78.14%	76.46%	75.01%	NA	NA

NOTE: Numbers captured on a GeForce GTX980

PERFORMANCE #2



	Grid Resolution				
	128 ³	256 ³	384 ³	512 ³	1,024 ³
	Full Grid				
Num. Bricks	256	2048	6,912	16,384	131,072
Memory (MB)	80	640	2,160	5,120	40,960
Simulation	2.29ms	19.04ms	64.71ms	NA	NA
	Sparse Grid				
	36	146	183	266	443
	11.25	45.63	57.19	83.13	138.44
Simulation	0.41ms	1.78ms	2.67ms	2.94ms	5.99ms
Scaling Sim.	78.14%	76.46%	75.01%	NA	NA

NOTE: Numbers captured on a GeForce GTX980

SCALING

- Ratio (in time) of 1 Brick = $\frac{\text{Time}\{\text{Full}\}}{\text{Time}\{\text{Sparse}\}}$
- ~75% across grid resolutions.

	Grid Resolution				
	128 ³	256 ³	384 ³	512 ³	1,024 ³
Scaling Sim.	78.14%	76.46%	75.01%	NA	NA

SUMMARY

- ▶ Let's see more fluid in games.
 - ▶ Fluid is not box shaped!
 - ▶ One volume is better than many small.
 - ▶ Un/Compressed storage a viable fallback.
 - ▶ CPU read-backs are useful if done right!
 - ▶ VTRs great for fluid simulation.
-
- ▶ Other latency resistant algorithms with tiled resources?



THANK YOU

JOIN THE CONVERSATION

#GTC15



▶ **ALEX DUNN** - ADUNN@NVIDIA.COM

▶ **TWITTER:** @ALEXWDUNN